

PATENT

RECEIVED
CENTRAL FAX CENTER
JUL 03 2008

Atty Docket No.: 200308654-1
App. Ser. No.: 10/705,932

IN THE CLAIMS:

Please find below a listing of all of the pending claims. The statuses of the claims are set forth in parentheses.

1. (Currently Amended) A method for identifying samples to determine search results for a query in a peer-to-peer system, the method comprising:

receiving a query at a destination node;

receiving samples from a first set of nodes proximally located to the destination node in an overlay network for the peer-to-peer system, the samples associated with information stored at the proximally located nodes;

identifying and selecting, based on the samples received from the first set of nodes, a first node of the first set of nodes likely storing information associated with objects stored in the peer-to-peer system that are relevant to the query to determine search results for the query;

storing forwarding the query to information for the identified first node of the first set of nodes likely storing information associated with objects stored in the peer-to-peer system that are relevant to the query;

generating semantic vectors for the objects stored in the peer-to-peer system;

hashing each of the semantic vectors to generate keys identifying locations in the overlay network to store key-value pairs for the objects, wherein the keys are the semantic vectors for the objects and the values include at least one of the objects and addresses for the objects; and

PATENT**Atty Docket No.: 200308654-1
App. Scr. No.: 10/705,932**

storing the key-value pairs at the nodes associated with the identified locations in the overlay network wherein the stored key-value pairs associated with similar semantic vectors are proximally located in the overlay network.

2. (Original) The method of claim 1, further comprising:

comparing the query to information stored in the first node; wherein the information stored in the first node is associated with objects stored in the peer-to-peer network; and generating search results including information stored in the first node associated with objects relevant to the query based on the comparison of the query to the information stored in the first node.

3. (Original) The method of claim 2, further comprising:

determining whether a quit threshold has been reached; transmitting the search results to an initiator of the query in response to the quit threshold being reached; and

performing the following steps in response to the quit threshold not being reached:
identifying a second node likely storing information associated with objects stored in the peer-to-peer network that are relevant to the query based on samples received from a second set of nodes including the second node, wherein the second set of nodes are nodes proximally located to the first node in the overlay network; and adding information stored in the second node to the search results; the added information being associated with objects that are relevant to the query.

PATENT**Atty Docket No.: 200308654-1
App. Ser. No.: 10/705,932**

4. (Previously Presented) The method of claim 3, wherein the quit threshold is associated with at least one of hops in the overlay network and whether the search results are improved by adding information to the search results from the second node.

5. (Cancelled).

6. (Previously Presented) The method of claim 1, further comprising:
generating the samples for the first set of nodes as a function of at least one of key-value pairs stored at each of the first set of nodes.

7. (Original) The method of claim 6, wherein generating the samples comprises:
generating a destination node semantic vector representative of objects associated with at least one of key-value pairs stored at the destination node and recent queries executed by the destination node;
generating a list of key-value pairs for each node of the first set of nodes, wherein each list includes key-value pairs associated with objects having semantics similar to the destination node semantic vector.

8. (Original) The method of claim 7, wherein identifying, based on the samples received from the first set of nodes, a first node of the first set of nodes likely storing information associated with objects stored in the peer-to-peer network that are relevant to the query comprises:

generating a semantic vector for each of the samples for the first set of nodes;

PATENT**Atty Docket No.: 200308654-1
App. Ser. No.: 10/705,932**

comparing the destination node semantic vector to each of the semantic vectors for the first set of nodes; and

identifying one of the semantic vectors for the first set of nodes closest to the destination node semantic vector.

9. (Previously Presented) The method of claim 1, further comprising:

identifying lower elements for the semantic vectors;

generating planes in the overlay network associated with the lower elements; and

performing the steps of claim 1 for each of the planes.

10. (Previously Presented) The method of claim 1, further comprising:

storing indices of key-value pairs at the nodes;

replicating an index for a second node in the first node, wherein the second node is proximally located to the first node in the overlay network; and

identifying key-value pairs from the replicated index that are relevant to the query.

11. (Previously Presented) The method of claim 1, further comprising:

storing indices of key-value pairs at the nodes;

in the first node, replicating indices for a plurality of nodes in a region in the overlay network including the first node; and

identifying key-value pairs from the replicated indices that are relevant to the query.

PATENT**Atty Docket No.: 200308654-1
App. Ser. No.: 10/705,932**

12. (Original) The method of claim 2, wherein the first set of nodes are neighbor nodes to the destination node in the overlay network.

13. (Original) The method of claim 3, wherein the second set of nodes are neighbor nodes to the first node in the overlay network.

14. (Currently Amended) An apparatus for identifying samples to determine search results for a query in a peer-to-peer system, the apparatus comprising:

means for receiving a query at a destination node;

means for receiving samples from a first set of nodes proximally located to the destination node in an overlay network for the peer-to-peer system, the samples associated with information stored at the proximally located nodes;

a storage device to store the received samples;

means for identifying and selecting, based on the samples received from the first set of nodes, a first node of the first set of nodes likely storing information associated with objects stored in the peer-to-peer system that are relevant to the query to determine search results for the query;

means for storing forwarding information for the query to the identified first node of the first set of nodes likely storing information associated with objects stored in the peer-to-peer system that are relevant to the query;

means for generating semantic vectors for the objects stored in the peer-to-peer system;

PATENT

Atty Docket No.: 200308654-1
App. Ser. No.: 10/705,932

means for hashing each of the semantic vectors to generate keys identifying locations in the overlay network to store key-value pairs for the objects, wherein the keys are the semantic vectors for the objects and the values include at least one of the objects and addresses for the objects; and

means for storing the key-value pairs at the nodes associated with the identified locations in the overlay network wherein the stored key-value pairs associated with similar semantic vectors are proximally located in the overlay network.

15. (Original) The apparatus of claim 14, further comprising:

means for comparing the query to information stored in the first node; wherein the information stored in the first node is associated with objects stored in the peer-to-peer network; and

means for generating search results including information stored in the first node associated with objects relevant to the query based on the comparison of the query to the information stored in the first node.

16. (Original) The apparatus of claim 15, further comprising:

means for determining whether a quit threshold has been reached;

means for transmitting the search results to an initiator of the query in response to the quit threshold being reached; and

means for performing the following functions in response to the quit threshold not being reached:

PATENT**Atty Docket No.: 200308654-1
App. Ser. No.: 10/705,932**

identifying a second node likely storing information associated with objects stored in the peer-to-peer network that are relevant to the query based on samples received from a second set of nodes including the second node, wherein the second set of nodes are nodes proximally located to the first node in the overlay network; and adding information stored in the second node to the search results; the added information being associated with objects stored in the peer-to-peer system that are relevant to the query.

17. (Previously Presented) The apparatus of claim 16, wherein the quit threshold is associated with at least one of hops in the overlay network and whether the search results are improved by adding information to the search results from the second node.

18. (Currently Amended) A computer readable storage device on which is embedded a program, the program performing a method, the method comprising:

receiving a query at a destination node;
receiving samples from a first set of nodes proximally located to the destination node in an overlay network for the peer-to-peer system, the samples associated with information stored at the proximally located nodes;
identifying and selecting, based on the samples received from the first set of nodes, a first node of the first set of nodes likely storing information associated with objects stored in the peer-to-peer system that are relevant to the query;

generating semantic vectors for the objects stored in the peer-to-peer system;

PATENT**Atty Docket No.: 200308654-1
App. Ser. No.: 10/705,932**

hashing each of the semantic vectors to generate keys identifying locations in the overlay network to store key-value pairs for the objects, wherein the keys are the semantic vectors for the objects and the values include at least one of the objects and addresses for the objects; and

storing the key-value pairs at the nodes associated with the identified locations in the overlay network wherein the stored key-value pairs associated with similar semantic vectors are proximally located in the overlay network.

19. (Previously Presented) The computer readable storage device of claim 18, wherein the method further comprises:

comparing the query to information stored in the first node; wherein the information stored in the first node is associated with objects stored in the peer-to-peer network; and

generating search results including information stored in the first node associated with objects relevant to the query based on the comparison of the query to the information stored in the first node.

20. (Previously Presented) The computer readable storage device of claim 19, wherein the method further comprises:

determining whether a quit threshold has been reached;

transmitting the search results to an initiator of the query in response to the quit threshold being reached; and

performing the following steps in response to the quit threshold not being reached:

PATENTAtty Docket No.: 200308654-1
App. Ser. No.: 10/705,932

identifying a second node likely storing information associated with objects stored in the peer-to-peer network that are relevant to the query based on samples received from a second set of nodes including the second node, wherein the second set of nodes are nodes proximally located to the first node in the overlay network; and adding information stored in the second node to the search results; the added information being associated with objects stored in the peer-to-peer system that are relevant to the query.

21. (Previously Presented) The computer readable storage device of claim 20, wherein the quit threshold is associated with at least one of hops in the overlay network and whether the search results are improved by adding information to the search results from the second node.

22. (Currently Amended) A peer-to-peer system comprising:

a plurality of nodes in the system operating as a search engine configured to execute a query received by the search engine, each of the plurality of nodes comprising a storage device to store information;

an overlay network implemented by the plurality of nodes;

a plurality of indices stored at the plurality of nodes, each index including at least one semantic vector for an object, and

each of the semantic vectors is being hashed to generate keys identifying locations in the overlay network to store key-value pairs for the objects,

wherein the keys are the semantic vectors for the objects and the values include at least one of the objects and addresses for the objects,

PATENT**Atty Docket No.: 200308654-1
App. Ser. No.: 10/705,932**

wherein the key-value pairs are stored at the nodes associated with the identified locations in the overlay network, and the stored key-value pairs that are associated with similar semantic vectors are proximally located in the overlay network;

wherein a first node in the search engine is configured to receive samples from the nodes proximally located to the first node in the overlay network, the first node utilizing the samples to identify and select an index of one of the other nodes to search in response to receiving the query.

23. (Cancelled).

24. (Currently Amended) The system according to claim 22 23, wherein the first node is located in a region in the overlay network and the first node is configured to store indices from nodes in the region, wherein the first node is operable to search a plurality of indices likely to include information relevant to the query without forwarding the query to other nodes in the region.